

AFRL-IF-RS-TM-2007-9
Final Technical Memorandum
April 2007



LEARNING ROBUSTLY FROM INSTRUCTIONS

Carnegie Mellon University

APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

STINFO COPY

AIR FORCE RESEARCH LABORATORY
INFORMATION DIRECTORATE
ROME RESEARCH SITE
ROME, NEW YORK

NOTICE AND SIGNATURE PAGE

Using Government drawings, specifications, or other data included in this document for any purpose other than Government procurement does not in any way obligate the U.S. Government. The fact that the Government formulated or supplied the drawings, specifications, or other data does not license the holder or any other person or corporation; or convey any rights or permission to manufacture, use, or sell any patented invention that may relate to them.

This report was cleared for public release by the Air Force Research Laboratory Rome Research Site Public Affairs Office and is available to the general public, including foreign nationals. Copies may be obtained from the Defense Technical Information Center (DTIC) (<http://www.dtic.mil>).

AFRL-IF-RS-TM-2007-9 HAS BEEN REVIEWED AND IS APPROVED FOR PUBLICATION IN ACCORDANCE WITH ASSIGNED DISTRIBUTION STATEMENT.

FOR THE DIRECTOR:

/s/

/s/

PETER J. ROCCI
Work Unit Manager

JOSEPH CAMERA, Chief
Information & Intelligence Exploitation Division
Information Directorate

This report is published in the interest of scientific and technical information exchange, and its publication does not constitute the Government's approval or disapproval of its ideas or findings.

REPORT DOCUMENTATION PAGE				<i>Form Approved</i> OMB No. 0704-0188	
<small>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Service, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington, DC 20503.</small>					
PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.					
1. REPORT DATE (DD-MM-YYYY) APR 2007		2. REPORT TYPE Final		3. DATES COVERED (From - To) Nov 04 – Nov 06	
4. TITLE AND SUBTITLE LEARNING ROBUSTLY FROM INSTRUCTIONS				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA8750-05-2-0014	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) John R. Anderson				5d. PROJECT NUMBER COGV	
				5e. TASK NUMBER 00	
				5f. WORK UNIT NUMBER 07	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University 5000 Forbes Ave. Pittsburgh PA 15213-3815				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFRL/IFED 525 Brooks Rd Rome NY 13441-4505				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSORING/MONITORING AGENCY REPORT NUMBER AFRL-IF-RS-TM-2007-9	
12. DISTRIBUTION AVAILABILITY STATEMENT APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED. PA# 07-198					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT The goal of the project was to perform a proof of potential by applying ACT-R, a modular architecture, to a challenging learning task. This task involved learning what is traditionally taught as the solution of linear equations in American high schools. This involved giving the system the abilities that a prepared student entering Algebra 1 should have. This included the abilities to perform basic arithmetic and to parse arithmetic Expressions, giving the system a representation of the instructions that appear in a standard algebra Textbook, and having the system learn by feedback on its solution efforts how to solve the class of problems that appear in the textbook.					
15. SUBJECT TERMS Cognitive Architectures, knowledge encoding, automated learning					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT UL	18. NUMBER OF PAGES 9	19a. NAME OF RESPONSIBLE PERSON Peter J. Rocci
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U			19b. TELEPHONE NUMBER (Include area code)

Table of Contents

1.0 Overview	1
2.0 Description of Work	2
3.0 Evaluation and Discussion	5
4.0 Technology	5
5.0 Publication	5

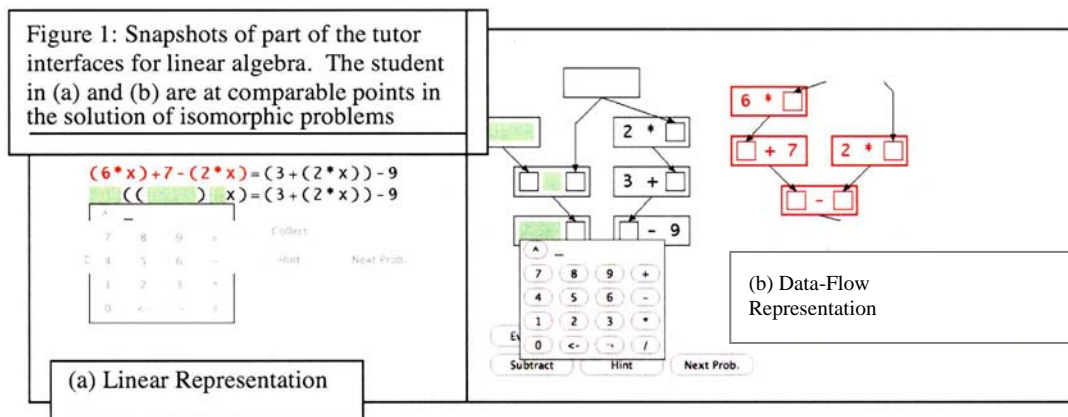
1.0 Overview

ACT-R (Adaptive Control of Thought – Rational) is a modular architecture that allows one to develop and integrate perceptual, motor, and more deliberative components. Because these components can be identified with specific brain areas, neural imaging can guide the development of the architecture. Its hybrid architecture integrates reflexive subsymbolic computations with deliberative symbolic computations. Its learning mechanisms reinforce this symbolic-subsymbolic duality in domains such as learning from instruction. ACT-R can take instruction and use it to guide its behavior symbolically. While initial instruction interpretation is very slow, it can compile new rules to efficiently apply this knowledge (symbolic learning). However, it can also learn how well these rules work in different contexts (subsymbolic learning) and how to select most likely paths. The goal of the project was to perform a proof of potential by applying ACT-R to a challenging learning task that we believe is just tractable. This task involved learning what is traditionally taught as the solution of linear equations in American high schools.

This involved:

1. Giving the system the abilities that a prepared student entering Algebra 1 should have. These include the abilities to perform basic arithmetic and to parse arithmetic expressions.
2. Giving the system a representation of the instructions that appear in a standard algebra textbook.
3. Having the system learn by feedback on its solution efforts how to solve the class of problems that appear in the textbook.

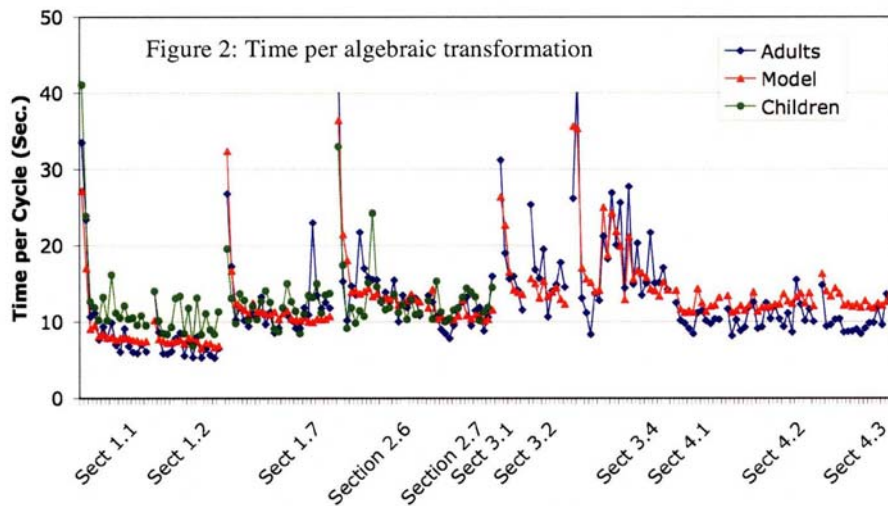
While it would be no challenge to build a system that solved linear equations or perhaps even to build a system that learned to do problems from a full specification of what it has to do, it would be a significant accomplishment to learn this on the same terms that people do. Learning algebra corresponds to an interesting phylogenetic transition. While there are now clever studies that have taught higher apes considerable arithmetic and perceptual skills (significant fractions of 1 above), algebra has a level of abstraction that makes it a uniquely human competence. So, while it is hardly the height of human intelligence, algebra is one of the more tractable reflections of human intelligence. As evidence that the ACT-R model was acquiring the competence in human terms we performed a functional MRI imaging study of the acquisition of this skill and tested predictions of the model for this imaging data as well as the behavioral data.



2.0 Description of Work

In the first year we created a computer tutor that covers the full curriculum on linear expressions and their solutions as contained in the first four chapters of the classic algebra textbook of Foerster. We also developed an isomorphic representation of this system in which equations are represented by data-flow graphs. It serves as a basis for teaching high-school algebra to college students in a form they do not recognize and have to learn over again. Figure 1 shows the representation of the same equation in the two systems. Both children and adults went through subsets of the curriculum within their respective systems. Children generally did somewhat worse than adults, but their errors were largely a matter of slips. If a student makes a slip and produces a wrong result it can be difficult to diagnose where the error was made and most of the extra time and errors for children involved recovery attempts from these initial slips. It turns out there is one measure on which children and adults do not differ. This is the time to perform a transformation of the equation. Because of the design of the interfaces, the number of motor actions to perform a transformation does not vary with form of the algebraic representation (linear or data flow). While the number of transformations will vary with the complexity of an equation, the time to perform a single transformation is also constant. Figure 2 shows the time per transformation across problems in the curriculum and the predictions of an ACT-R model for this task. The most striking aspect of the data comes from the large spikes in the times where new instructions are being introduced. This reflects the cost of the initial interpretation of these instructions. The ACT-R model parsed the instruction delivered by the tutor and encoded it in a declarative form. The knowledge is encoded in the form of operators that the agent can perform, the preconditions of these operators, and their post conditions. Besides encoding mathematical knowledge, this operator encoding has been used to model knowledge in domains as widely varied as simple dual-task experiments (Anderson et al., 2005) and operating a flight management system on a Boeing 777 (Taatgen et al., 2006). An example of such an operator might be how to distribute multiplication over addition in an algebraic expression, what the

preconditions are for this operation, and what the consequences are. Encoding of mathematical knowledge into a rich set of operators for a domain provides the student with the ability to reason about a wide range of representations of mathematical relationships and their implications. However, knowledge in this highly flexible form is expensive to use in terms of working memory demands. As the student experiences repeated patterns of this knowledge use, new procedures develop for efficiently using this knowledge without retrieval and application of the original declarative encoding of the instructions.



As is apparent from Figure 2, the theory does more than just predict these general behavioral patterns; it predicts the actual observed data. Still, such behavioral data are too coarse-brained to substantiate fully the detailed modeling assumptions. Therefore, in the second year we turned to fMRI brain imaging to obtain converging evidence about the specific operations in the theory. The system works well in an fMRI scanner and we ran college students through the curriculum. Figure 3 shows the results from four important brain regions. It shows the results for simple and complex equations early and late in learning. The length of the BOLD curves reflects the duration of the equation solving while the area under the curves reflect the amount of energy being expended in that region of the brain. Each of the four regions shows a distinct pattern:

- (a) The motor region reflects the number of actions to execute the operation, which is greater for more complex equations but which does not change with practice. Therefore, there is greater area under the curve in the case of the more complex equations but no different between early versus late in learning.
- (b) The parietal region reflects the amount of time spent in encoding operations to represent the equation and its transformations. The area under these curves is greater for more complex equations and reduces with practice because certain encoding operations become skipped.

(c) The prefrontal region reflects retrieval of instructions and arithmetic facts. It shows a similar learning pattern to the parietal except that the effect of practice is even more dramatic because all instruction retrieval has dropped out.

(d) The anterior cingulate reflects the decision making being made at different choice points in the solution of the equations. This largely does not change with practice. However, it turns out that there is actually an increase late in practice with complex equations. As part of the learning, students learn more complex transformations that complicate the decisions about the signs of terms.

Therefore, reflecting their greater knowledge they actually have to make more decisions later. The solid lines in the figures represent the predictions of the ACT-R model and the dotted lines connect the actual data points. It can be seen that the correspondence is pretty close.

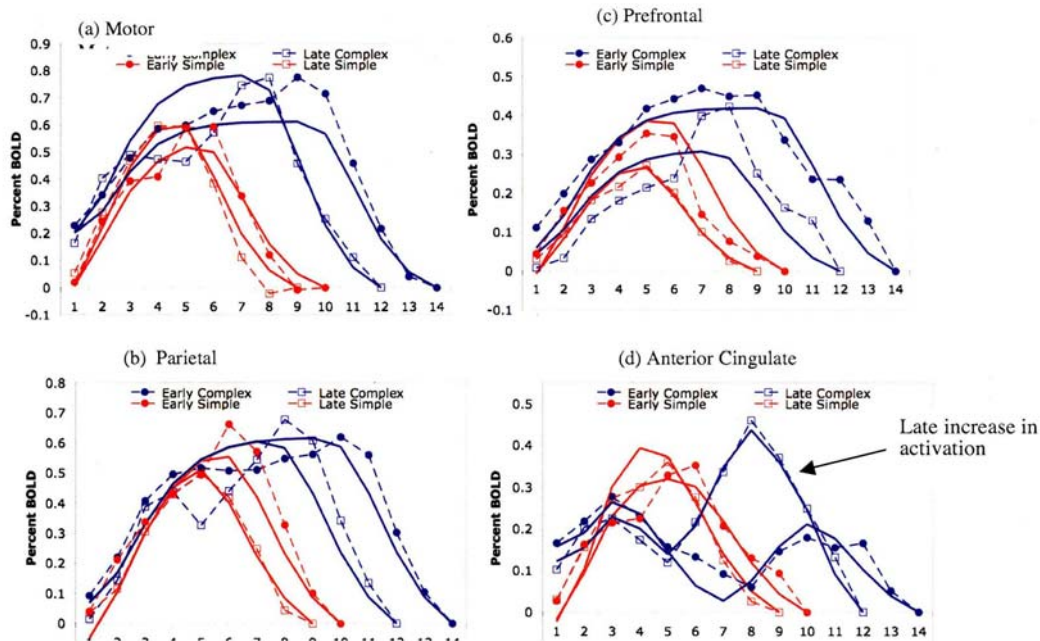


Figure 3: Preliminary data (points connected by dotted lines) and model predictions (solid lines) for fMRI results from algebra tutor

3.0 Evaluation and Discussion

As Figures 2 & 3 display, the model did a good job in achieving a correspondence with the detailed data obtained from students. Probably the two outstanding issues about the approach are whether the models detailed assumptions are correct about instruction representation and whether the general approach would extend to other domains. With respect to the first question the best way to answer this would be to perform manipulations of the instructional input and see what the results were. We were able to perform such a study in the second year of the proposal but we did not have the time to extend the model to this data. We are doing this as part of another project funded by DARPA. With respect to the second question we have been looking at applying the same methodology operating a flight management system on a Boeing 777 as part of another grant.

4.0 Technology

The experimental system and ACT-R model as developed for this project is available at the following web site:

<http://act-r.psy.cmu.edu/models/>

5.0 Publication

The forthcoming book includes a major report on this research:
Anderson, J. R. (2007) How Can the Human Mind Occur in the Physical Universe? New York: Oxford University Press.